

A census of zeta functions of quartic K3 surfaces over \mathbb{F}_2

Kiran S. Kedlaya and Andrew V. Sutherland

ABSTRACT

We compute the complete set of candidates for the zeta function of a K3 surface over \mathbb{F}_2 consistent with the Weil conjectures, as well as the complete set of zeta functions of smooth quartic surfaces over \mathbb{F}_2 . These sets differ substantially, but we do identify natural subsets which coincide. This gives some numerical evidence towards a Honda-Tate theorem for transcendental zeta functions of K3 surfaces; such a result would refine a recent theorem of Taelman, in which one must allow an uncontrolled base field extension.

1. Introduction and results

For X an algebraic variety over a finite field \mathbb{F}_q of characteristic p , the *zeta function* of X is the power series

$$\zeta(X, T) = \exp \left(\sum_{n=1}^{\infty} \#X(\mathbb{F}_{q^n}) \frac{T^n}{n} \right).$$

A number of basic properties of $\zeta(X, T)$ are controlled by the now-proved *Weil conjectures* (see for example [21]); for example, $\zeta(X, T)$ always represents a rational function in $\mathbb{Q}(T)$.

Given a class of varieties over a particular field, it is natural to pose the *inverse problem* asking which zeta functions consistent with the Weil conjectures actually occur. For abelian varieties of a given dimension g over \mathbb{F}_q , this question is resolved by celebrated theorems of Tate [25] and Honda [12]: all such zeta functions occur provided that if $q \neq p$, one adds an extra condition on the factorizations over \mathbb{Q} and \mathbb{Q}_p . (This condition always holds in the ordinary case; see [10, §4] for a concise statement of the condition and [20] for a thorough exposition.) By contrast, for curves of a given genus g , there are many additional constraints (the inequalities $\#X(\mathbb{F}_q) \geq 0$ and $\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n})$, for example), and even the maximum value of $\#X(\mathbb{F}_q)$ is unknown in most cases (see [19] for some results).

In this paper, we make a numerical investigation of the inverse problem for zeta functions of K3 surfaces over \mathbb{F}_2 . Recall that a *K3 surface* over \mathbb{F}_q is a smooth, simply connected[†] projective surface with trivial canonical bundle. The geometry of K3 surfaces is in many ways analogous to that of elliptic curves (they are Calabi-Yau varieties of dimensions 2 and 1, respectively). However, one key difference is that K3 surfaces cannot be uniformly described using a single geometric construction. Instead, an infinite number of distinct constructions are required; we will focus mainly on the case of smooth quartic (degree 4) surfaces in \mathbb{P}^3 .

In the case of a K3 surface, one reads off from the Weil conjectures (plus properties of crystalline cohomology) the following constraints; see [24, Theorem 1] for references and a sharper statement in case q is not prime (in the same vein as the Honda-Tate theorem).

2000 *Mathematics Subject Classification* 11M38, 14J28.

Kedlaya was supported by NSF grant DMS-1501214, UCSD (Warschawski chair), and a Guggenheim Fellowship. Sutherland was supported by NSF grants DMS-1115455 and DMS-1522526.

[†]Here we mean that the surface admits no nontrivial connected finite étale covers. This condition is needed only to eliminate abelian surfaces.

THEOREM 1. *Let X be a K3 surface over \mathbb{F}_q . Then $\zeta(X, T)$ has the form*

$$\frac{1}{(1-T)(1-qT)(1-q^2T)q^{-1}L(qT)}$$

for some polynomial $L(T) \in \mathbb{Z}[T]$ of degree 21 with $L(0) = q$ having all roots on the unit circle.

One also has the following consequence of the Artin-Tate formula [7]; see Theorem 8 for more discussion.

THEOREM 2. *With notation as in Theorem 1, write $L(T) = (1-T)^r L_1(T)$ with $L_1(1) \neq 0$. Then $L_1(-1)$ is a perfect square (possibly 0).*

With these preliminaries in hand, we describe our computational results concerning zeta functions of K3 surfaces over \mathbb{F}_2 ; the code used for these computations can be found in the repository <https://github.com/kedlaya/root-unitary>. Our first computational result is an enumeration of Weil polynomials based on a refinement of the search strategy described in [11]; see §2 for details.

COMPUTATION 3. *The following sets are computed.*

- (a) *The set of polynomials $L(T)$ satisfying the conditions of Theorem 1 for $q = 2$; it contains 2,971,182 elements.*
- (b) *The set of polynomials in (a) consistent with Theorem 2; it contains 2,195,801 elements.*
- (c) *The set of polynomials in (b) consistent with the inequalities $\#X(\mathbb{F}_q) \geq 0$ and $\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n})$ (it suffices to impose the second condition for $(mn, n) \in \{(2, 1), (3, 1), (4, 2)\}$); it contains 1,672,565 elements.*

Our second computational result is a lower bound for the inverse problem obtained by enumerating smooth quartic surfaces X/\mathbb{F}_2 and computing $\zeta(X, T)$ directly by counting points in $X(\mathbb{F}_{2^n})$; see §3 for details.

COMPUTATION 4. *The following sets are computed.*

- (a) *The set of $\mathrm{PGL}_4(\mathbb{F}_2)$ -equivalence classes of smooth quartic surfaces over \mathbb{F}_2 ; it contains 528,257 elements.*
- (b) *The set of zeta functions of the surfaces in (a); it contains 52,755 elements and is a subset of the set found in Computation 3(c).*

With regard to (a), note that distinct PGL_4 -equivalence classes with the same zeta function may in fact give rise to isomorphic K3 surfaces: within the Néron-Severi lattice of a single K3 surface, the ample cone may contain multiple inequivalent divisors of degree 4. However, this cannot occur for $r = 1$ (i.e., when $L'(1) \neq 0$). With regard to (b), note that in loose analogy with Tate's theorem that isogenous abelian varieties have the same zeta function [25], a theorem of Lieblich and Olsson [14, Theorem 1.2] and Huybrechts [13, Proposition 4.6] states that K3 surfaces which are *derived equivalent* (or *Fourier-Mukai equivalent*) have the same zeta function.

Computation 3 provides a rich data set for investigating questions about zeta functions of K3 surfaces; for example, all possible values $1, \dots, 10, \infty$ for the height of a K3 surface are realized by smooth quartics over \mathbb{F}_2 . This said, the meaning of Computation 4 for the inverse

problem for K3 surfaces over \mathbb{F}_2 is unclear. However, it does yield some evidence towards a weaker form of the inverse problem suggested by Taelman [24]. For $L(T)$ as in Theorem 1, factor $L(T) = \prod_i (1 - \alpha_i T)$ over \mathbb{C} and define the *algebraic part* and *transcendental part*

$$L_{\text{alg}}(T) := \prod_{i: \alpha_i \in \mu_\infty} (1 - \alpha_i T), \quad L_{\text{trc}}(T) := \prod_{i: \alpha_i \notin \mu_\infty} (1 - \alpha_i T),$$

where μ_∞ denotes the group of roots of unity in \mathbb{C}^\times . (By Computation 3, for $q = 2$ there are 73,617 possible values of $L_{\text{trc}}(T)$, whether or not we add conditions (b) and (c).) Then one can pose the inverse problem for $L_{\text{trc}}(T)$ in place of $L(T)$, and in this case one has the following (conditional) partial solution [24, Theorem 2].

THEOREM 5. *Assume that all K3 surfaces over finite extensions of \mathbb{Q}_p have potential semistable reduction[†] in the sense of [24, Definition 1]. Let $L_{\text{trc}}(T) = \prod_i (1 - \alpha_i T)$ be a polynomial arising from some $L(T)$ as in Theorem 1; if q is not prime, impose also the additional restrictions given in [24, Theorem 1]. Then for some positive integer n , the polynomial $\prod_i (1 - \alpha_i^n T)$ occurs as the transcendental part for some K3 surface over \mathbb{F}_{q^n} .*

The proof of [24] gives little insight as to whether the conclusion should necessarily hold with $n = 1$. However, we can use our preceding computations to give a statement in this direction.

COMPUTATION 6. *The following sets are computed.*

- (a) *The subset of Computation 3(a) for which $L_{\text{alg}}(T) = 1 + T$, $L_{\text{trc}}(1) = 2$, $L_{\text{trc}}(-1) > 2$; it contains 1,995 elements. (Adding the conditions of Computation 3(c) does not change this answer.)*
- (b) *The corresponding subset of Computation 4(b); it contains the same 1,995 elements.*

The conditions imposed in Computation 6 were chosen to partially (but not completely) eliminate the possibility that $L(T)$ arises from a K3 surface other than a smooth quartic by accounting for the Artin-Tate formula. It would be natural to continue the analysis by considering other families of K3 surfaces; however, the Artin-Tate formula makes it difficult to produce enough examples to establish that Theorem 5 always holds for $q = 2$ with $n = 1$. See §4 for elucidation of this point.

In another direction, one may hope to make similar calculations for $q > 2$, but this poses significant technical challenges. Again, see §4 for further details.

2. Tabulation of Weil polynomials

Our tabulation of Weil polynomials broadly follows the search strategy described in [11, §5]; it is similar in spirit to the tabulation of number fields of prescribed signature, as in the work of Malle [18] and Voight [26]. We briefly recall the strategy, indicate some new refinements which make it feasible to conduct much larger searches than previously possible, and discuss the computations performed.

[†]This condition, which would be guaranteed if we were in equal characteristic 0 and is likely to hold in general, is only known for K3 surfaces with a polarization of small degree relative to p . See [16, §2] for more discussion.

2.1. The search strategy

This strategy attacks the problem of tabulating (not necessarily monic) integer polynomials

$$P(T) = a_n T^n + \cdots + a_0$$

of a fixed degree whose roots lie on the unit circle, with a_0 fixed and a_1, \dots, a_n constrained to lie in certain congruence classes (which could be a singleton set if the modulus is 0, or all integers if the modulus is 1). All roots other than ± 1 occur in conjugate pairs; hence if n is odd, then one of ± 1 occurs with odd multiplicity, and we may reduce to the case where $n = 2m$ is even and $a_i = a_{n-i}$ for all i . (For example, for K3 surfaces we have $n = 21$, so we first find reciprocal polynomials of degree 20, then multiply each one by $1 + T$ and $1 - T$ to generate the desired list.) We may then write

$$P(T) = T^m Q(T + 1/T), \quad Q(T) = b_m T^m + \cdots + b_0$$

for some integer polynomial $Q(T)$ with roots in the interval $[-2, 2]$, with b_m fixed and b_0, \dots, b_{m-1} constrained to congruence classes U_0, \dots, U_{m-1} .

To enumerate the set S of polynomials $Q(T)$ we compute a tree with levels $0, \dots, m$ in which each node (b_m, \dots, b_0) at level m represents a polynomial $Q(T) = b_m T^m + \cdots + b_0$ in S , and nodes at level $i < m$ are labeled by tuples (b_m, \dots, b_{m-i}) that are prefixes of their children (b_m, \dots, b_{m-i-1}) . By Rolle's theorem, a necessary (but not sufficient) condition for the node (b_m, \dots, b_k) to have a descendant (b_m, \dots, b_0) at level m is that the polynomial

$$\sum_{j=0}^{m-k} \binom{k+j}{j} b_{k+j} T^j$$

has all its roots in $[-2, 2]$. A general description of the algorithm appears below; a particular implementation is described in the next subsection.

ALGORITHM 7. Given b_m and congruence classes U_0, \dots, U_{m-1} , enumerate the nodes of a tree with root (b_m) at level 0 in which each node at level $i > 0$ is an integer tuple (b_m, \dots, b_{m-i}) with parent (b_m, \dots, b_{m-i+1}) as follows.

- (a) Given a node (b_m, \dots, b_{m-i}) , check whether the polynomial

$$R(T) := \sum_{j=0}^i \binom{m-i+j}{j} b_{m-i+j} T^j$$

has all its roots in $[-2, 2]$.

- (b) If the test in (a) passes and $i = m$, add $b_m T^m + \cdots + b_0$ to a list of return values.
(c) If the test in (a) passes and $i < m$, compute an interval I with the following property: for any values b_{m-i-1}, \dots, b_0 such that $b_m T^m + \cdots + b_0$ has all roots in $[-2, 2]$, one has $b_{m-i-1} \in I$. Then take the children of this node to be the tuples (b_m, \dots, b_{m-i-1}) with $b_{m-i-1} \in U_{m-i-1} \cap I$ (this intersection may be the empty set).

2.2. Implementation

An implementation[†] of the aforementioned search strategy (available for download) is described in [11]. The implementation we use here differs from the prior one in several theoretical and practical aspects, which we now discuss.

[†]While preparing this paper, we discovered a minor bug in the implementation accompanying [11]. However, we corrected this bug in the current implementation and reconfirmed all of the computational results.

- In [11], the interval I in Algorithm 7(c) is constructed using linear and quadratic inequalities on the power sums s_1, \dots, s_{i+1} , as computed from b_m, \dots, b_{m-i-1} via the Newton identities; note that m appears explicitly in the identities, so these inequalities typically carry more information than simply requiring that the one-step extension conform to Rolle’s theorem. See [11, §5] for the precise list of inequalities used.

We add a new condition which is linear in the b_j rather than the s_j : if $Q(T)$ has all roots in $[-2, 2]$, then $b_m Q(2 + T)$ and $(-1)^m b_m Q(2 - T)$ have all coefficients nonnegative (consistent with Descartes’s rule of signs).

- In [11], the test in Algorithm 7(a) is conducted as follows. Let $S(T)$ be the polynomial obtained from $R(T)$ by removing all multiple roots and all factors of $T \pm 2$. Form a Sturm sequence S_0, S_1, \dots as follows: $S_0 := S$, $S_1 := S'$, and for $j > 1$, let S_i be a negative scalar multiple of the remainder of S_{j-2} modulo S_{j-1} , stopping just before appending the zero polynomial to the sequence. We then invoke Sturm’s theorem [2, Theorem 2.50]: the number of roots of $S(T)$ in $[-2, 2]$ is the difference between the numbers of sign changes in the sequences $S_0(-2), S_1(-2), \dots$ and $S_0(2), S_1(2), \dots$.

In the current implementation, we note that when testing R , we have the prior information that $R'(T)$ has all roots in $[-2, 2]$ and that $R(2)$, $R(-2)$ have the correct signs (thanks to the previous point). We thus need only test that R has all real roots; this eliminates some polynomial evaluations at ± 2 . We may also take $S = R$, since the fact that Sturm’s theorem remains valid as a count of roots without multiplicity; this avoids duplication of the Euclidean algorithm. We next observe that R has all real roots if and only if for $j = 1, \dots, i$, $\deg(S_j) = i - j$ and the leading coefficients of S_0 and S_j have the same sign; this allows for an early abort. Finally, we note that if the early abort happens due to a sign discrepancy (rather than a degree discrepancy) at S_j , then this discrepancy does not depend on b_k for $k \leq m - i - 1 + \max\{0, i - 2j + 2\}$; if $i - 2j + 2 > 0$, we may thus back up the tree traversal and discard all nodes below (b_m, \dots, b_{m-2j+2}) without losing any of the return values. (If $b_m > 0$ and nodes are traversed in lexicographic order, then one may replace $i - 2j + 2$ with $i - 2j + 1$ and b_{m-2j+2} with b_{m-2j+1} in the previous analysis.)

- In [11], the implementation consisted of an interpreted component in Sage [23] performing high-level user interaction and a compiled component in Cython [6] for mid-level computations. Some low-level computations, such as Sturm sequences, were farmed out to compiled components of the Sage library, notably PARI [22].[†]

In the current implementation, we incorporate a third component, written in C using the FLINT library [8]. This component absorbs most of the work of the Cython layer (which remains to provide wrappers around the C code) and completely supplants the use of PARI.

- In [11], parallelization via work-stealing is suggested but not implemented; we provide this in the current implementation. Given a pool of threads, we initially assign the entire search tree to one thread, then iterate the following steps until no active threads remain.
 - (i) In parallel, each thread which is active (i.e., has been assigned a subtree of the search tree) performs a depth-first search to find one polynomial in its remaining search space, going inactive if none exist.
 - (ii) In serial, each inactive thread solicits work from a randomly chosen active thread by removing a branch (as close to the root as possible) from the latter’s subtree.

[†]The PARI/GP project includes both the C library PARI and the interpreted GP language. Sage interfaces directly with the C library.

2.3. Computations

We now describe in detail some computational results obtained using this search strategy. The reported computations were carried out on a 24 core Intel Xeon X5690 3.47GHz machine with 192GB of memory. The parallel implementation was run using 512 threads; this provided an 8–10 \times speedup.

Computation 3 was completed in under 1 hour. In addition to the 1,485,591 polynomials of degree 20 that were found, the search tree found an additional 2,149,281,061 leaves at smaller depths; that is, the dead ends outnumber the solutions by a factor of nearly 1,500. This suggests that there may still be substantial room to optimize the choice of the intervals in Algorithm 7(b) (see below).

As an additional test of the implementation (which helped expose some bugs during development), we computed the set of monic polynomials $L(T) \in \mathbb{Z}[T]$ with degree in $\{1, 3, \dots, 21\}$ with all roots on the unit circle; it contains 78,670 elements. By Kronecker’s theorem, these polynomials factor as products of cyclotomic polynomials; it is thus easy to confirm this answer using an independent script that forms these products directly.

2.4. Possible refinements

The discussion above suggests that there remains substantial room for improvement in the choice of intervals in Algorithm 7(b). As noted in [11], similar searches in other contexts often make use of linear programming methods; we have not investigated this direction.

There is also a question of whether Sturm sequences are the optimal method for the test in Algorithm 7(a). For one, Sturm sequences involve multiprecision integers, in part due to the systemic appearance of certain large powers. In principle, these powers can be removed explicitly, thus reducing the computational complexity [4, Algorithm 3.3.1]; in practice, we find (in this particular setting) that the Gaussian content is substantially larger than predicted by general arguments, so we prefer to compute it explicitly. (One could try mixing the two approaches, but in [4, §3.3] it is suggested that this gives inferior results.)

More seriously, there is a question as to whether Sturm sequences are superior to root isolation methods based on the Budan-Fourier theorem [2, Theorem 2.35] (see also [2, §10.4]). We have chosen Sturm sequences in part for ease of implementation, but also because they are better suited to the task at hand. To wit, root isolation methods can easily generate certificates that guarantee the existence of certain real roots (using sign changes), but have more difficulty generating certificates that guarantee the failure of a polynomial to have all of its roots in an interval. By contrast, Sturm’s theorem provides certificates of the latter type easily using the early-abort mechanism described above. That said, it may be that a well-crafted strategy using root isolation (e.g., one which uses the positions of the roots of R' to help isolate the roots of R) would work better in the long run.

In order to get a fuller parallel speedup, some refinement of the parallelization mechanism is needed. For example, we currently only interrupt a process when it finds a solution or exhausts its search space; this is suboptimal in certain use cases where the search tree is very large but the number of solutions to be found is small.

3. Point counting

We computed the zeta function of every K3 surface over \mathbb{F}_2 that arises as a smooth quartic surface X in \mathbb{P}^3 by counting points on these surfaces over extension fields \mathbb{F}_{2^n} , with n ranging over a set of values sufficient to uniquely determine the zeta function $\zeta(X, T)$ given by Theorem 1 and Theorem 2 (up to $n = 19$ in the worst case). This computational problem naturally breaks down into tasks: (1) enumerate a complete set S of smooth surfaces defined

by homogeneous quartic polynomials $f \in \mathbb{F}_2[w, x, y, z]$, up to PGL_4 -equivalence; (2) compute $\#X(\mathbb{F}_{2^n})$ for $X \in S$ and suitable values of n .

3.1. Determining PGL_4 -orbits of homogeneous quartics

There are $\binom{7}{3} = 35$ homogeneous quartic monomials $w^a x^b y^c z^d$, one for each quadruple of nonnegative integers (a, b, c, d) with $a + b + c + d = 4$. If we order the quadruples (a, b, c, d) lexicographically, each homogeneous quartic $f \in \mathbb{F}_2[w, x, y, z]$ can be uniquely identified with a bit-vector $v := v(f) \in \mathbb{F}_2^{35}$ indexed by quadruples (a, b, c, d) for which $f = \sum v_{(a,b,c,d)} w^a x^b y^c z^d$; the bit-vector v can be conveniently encoded as an integer in $[0, 2^{35})$ and we order them accordingly (this is just the lexicographic order on $\{0, 1\}^{35}$).

The group $\mathrm{PGL}_4(\mathbb{F}_2)$ acts on the set of homogeneous quartics $f(w, x, y, z)$ via linear change of variables. In terms of the corresponding set $V := \mathbb{F}_2^{35}$ of vectors $v(f)$, each element of $\mathrm{PGL}_4(\mathbb{F}_2)$ corresponds to an invertible linear transformation of V that can be explicitly represented as an invertible 35×35 matrix. We may thus identify $\mathrm{PGL}_4(\mathbb{F}_2)$ with a subgroup G of $\mathrm{GL}_{35}(\mathbb{F}_2)$ of order $2^6(2^2 - 1)(2^3 - 1)(2^4 - 1) = 20,160$. As we are only interested in the quartic surfaces $f(w, x, y, z) = 0$ up to isomorphism, it suffices to consider the G -orbits of V , each of which may be uniquely represented by a lexicographically minimal v . The number of G -orbits can be computed via Burnside's lemma as

$$\#(V/G) = \frac{1}{\#G} \sum_g \#V^g = \frac{\#C}{\#G} \sum_C (\#\mathbb{F}_2)^{\dim_1(C)} = 1,732,564, \quad (3.1)$$

where the first sum is over group elements, the second sum is over conjugacy classes, and $\dim_1(C)$ denotes the dimension of the 1-eigenspace of the conjugacy class C . There are only 14 conjugacy classes in $\mathrm{PGL}_4(\mathbb{F}_2)$, so the second sum is trivial to compute.

To find lexicographically minimal representatives for each orbit we simply enumerated every orbit using a bitmap with 2^{35} entries; this took less than 2 days. We note that this brute-force approach is not feasible for finite fields larger than \mathbb{F}_2 . Indeed, determining a set of unique orbit representatives is already a nontrivial problem over \mathbb{F}_3 (the vector space \mathbb{F}_3^{35} contains $3^{35} \approx 2^{55.5}$ elements). However, determining the cardinality of this set via (3.1) is quite feasible for values of $q > 2$; for example, over \mathbb{F}_3 there are 4,127,971,480 orbits, and over \mathbb{F}_5 there are 100,304,466,278,983.

Having compiled a complete list of $\mathrm{PGL}_4(\mathbb{F}_2)$ -orbits of homogeneous quartics, we then want to restrict to those that define a K3 surface; this amounts to discarding orbits represented by a vector $v(f)$ for which the polynomial $f \in \mathbb{F}_2[w, x, y, z]$ is not irreducible, or for which the singular locus defined by the Jacobian matrix of f is nonempty (the latter implies the former but the former is often easier to check). These conditions are straightforward to apply, and we quickly find that 528,257 of the 1,732,564 orbits satisfy them; these constitute our set S of smooth plane quartic surfaces X/\mathbb{F}_2 in \mathbb{P}^3 .

3.2. Counting points on quartic surfaces over \mathbb{F}_2

Given a smooth quartic surface $X \in S$ defined by $f(w, x, y, z) = 0$, our basic strategy for computing $\#X(\mathbb{F}_{2^n})$ is elementary: iterate over pairs $(x_0, y_0) \in \mathbb{F}_{2^n}^2$ and for each pair determine the number of roots of the polynomial $g(w) := f(w, x_0, y_0, 1) = g(y) \in \mathbb{F}_{2^n}[w]$ that lie in \mathbb{F}_{2^n} (of course we also need to account for points with $z = 0$, but this reduces to the much easier problem of counting points on a curve in \mathbb{P}^2 and takes negligible time). To count the roots of $g(w)$ we use Zinoviev's formulas [27], which for low degree polynomials g over \mathbb{F}_{2^n} gives explicit $n \times n$ systems of linear equations over \mathbb{F}_2 whose solutions correspond to the roots of g , based on Berlekamp's algorithm for factoring polynomials over finite fields of small characteristic using linear algebra [3].

One might generically expect g to have degree 4, but in fact this is not the case. For all but 34 of the surfaces in S , the degree of the defining polynomial $f(w, x, y, z)$ in w is at most 3 (note that our lexicographic ordering minimizes the degree in w). In the typical case where $g(w)$ is a cubic, after making it monic and applying a linear change of variable we may assume $g(w) = w^3 + g_1w + g_0$. It is then feasible to precompute a lookup table T indexed by pairs $(g_0, g_1) \in \mathbb{F}_{2^n}^2$ whose entries record the number of roots of $w^3 + g_1w + g_0$ in \mathbb{F}_{2^n} . Each entry in T is an integer in $[0, 3]$ that can be encoded in 2 bits, thus the total size of T is 2^{2r+1} bits; even for $r = 19$, this is reasonably small (64 GB). The time to compute T is actually less than the time to instantiate $f(w, x_0, y_0, 1)$ at every pair $(x_0, y_0) \in \mathbb{F}_{2^n}^2$; we can accelerate this computation by enumerating the pairs (g_0, g_1) in an order that makes it convenient to compute the matrices appearing in Zinoviev’s formulas. Thus makes the computation of T worthwhile even for a single surface X , and we can reuse the same table T for every $X \in S$. The cost of point counting is then dominated by the time to evaluate $f(w, x_0, y_0, 1)$.

Given that we wish to count points on a fairly large set of surfaces (all of S for $n \leq 12$ and subsets $S_n \subseteq S$ for $n > 12$), rather than iterating over surfaces $f(w, x, y, z) = 0$ and counting points on each, which involves iterating pairs $(x_0, y_0) \in \mathbb{F}_{2^n}^2$ and evaluating $f(w, x_0, y_0, 1)$, we reverse the order of iteration and loop over pairs (x_0, y_0) and for each pair count the solutions to $f(w, x_0, y_0, 1) = 0$ over \mathbb{F}_{2^n} for every surface $f(w, x, y, z) = 0$ in our set, keeping a running total of points for each surface as we go. This allows us to instantiate the 35 homogeneous quartic monomials at $x = x_0, y = y_0, z = 1$ just once for each pair (x_0, y_0) , and then for each polynomial $f(w, x, y, z)$ compute $f(w, x_0, y_0, 1)$ as an \mathbb{F}_2 -linear combination of these, equivalently, as a sum of a subset of them, which is very fast.

This algorithm is trivially parallelizable (with linear speedup), and running on 32 cores it takes only two days to compute $\#X(\mathbb{F}_{2^n})$ for $1 \leq n \leq 12$ and all 528,257 surfaces $X \in S$. From these point counts, for each $X \in S$ we can write $L(T) = 1 + a_1T + \cdots + a_{21}T^{21}$ with a_1, \dots, a_{12} known. In most cases, the existence of a sign $\epsilon \in \{+, -\}$ such that $a_{21-i} = \epsilon a_i$ then determines $L(T)$ uniquely; the exceptions are the cases where $a_9 = a_{10} = 0$. For $n = 13, \dots, 21$, let S'_n be the subset of S consisting of exceptions for which $a_{21-n} \neq 0$ and $a_{20-n} = \cdots = a_{10} = 0$; let S_n be the subset of S_n for which both choices for ϵ give polynomials compliant with Theorem 1, Theorem 2, and Theorem 8 (which for degree 4 means that if $r = 1$, then $L_1(1)$ is a square). The sizes of these sets are listed in Table 1.

For $n = 13, \dots, 19$, we reran the previous algorithm to compute $\#X(\mathbb{F}_{2^n})$ for each $X \in S'_n$; the most time-consuming computation was for S'_{19} , which took about six days running on a machine with 32 Intel Xeon E5-2687Wv2 3.4GHz cores and 256GB of memory. (It would have been sufficient to consider $x \in S_n$, but the extra computations serve as a consistency check.) Note that the 1,876 tuples $(\#X(\mathbb{F}_2), \dots, \#X(\mathbb{F}_{2^{12}}))$ represented by $S_{13} \cup \cdots \cup S_{19}$ only give rise to 2,071 different zeta functions; that is, in the vast majority of these cases only one of the two sign choices is realized. This suggests that there may be further theoretical restrictions on zeta functions that we have not yet taken into account (e.g., interaction between the Newton polygon and the order of the Brauer group).

n	13	14	15	16	17	18	19	20	21
$\#S'_n$	38,225	16,555	8,281	3,608	2,011	857	283	0	96
$\#S_n$	17,795	7,315	3,611	1,435	1,016	470	125	0	0

TABLE 1. Cardinalities of sets S'_n, S_n

We should also mention an important low-level optimization to speed up arithmetic in \mathbb{F}_{2^n} that we used: the “carry-less multiplication” instruction PCLMULQDQ now available on Intel processors (since 2010) speeds up multiplication in \mathbb{F}_{2^n} quite dramatically (by a factor of up to 10 for the values of n that we used).

4. Further discussion

In light of the preceding computations, we resume the discussion from the introduction concerning the inverse problem for zeta functions of K3 surfaces.

As noted earlier, smooth quartics give rise to only one out of infinitely many algebraic families of K3 surfaces. This is due to the fact that, just as for abelian varieties, in order to represent the moduli problem one must consider *polarized* K3 surfaces. The *degree* of the polarization equals its self-pairing in the Néron-Severi lattice; each value of the degree corresponds to a single irreducible component of the moduli space of K3 surfaces. For the first few degrees, the generic polarized K3 surface of that degree can be described as follows. (There are also some nongeneric possibilities; see below for some discussion of the case of degree 4.)

- Degree 0: an elliptic K3 surface.
- Degree 2: a double cover of \mathbb{P}^2 (or a twist) branched over a smooth sextic (degree 6) curve.
- Degree 4: a smooth quartic in \mathbb{P}^3 .
- Degree 6: a smooth transversal intersection of a quadric and a cubic in \mathbb{P}^4 .
- Degree 8: a smooth transversal intersection of three quadrics in \mathbb{P}^5 .

For X a K3 surface over \mathbb{F}_q , define $L(T)$ as in Theorem 1 and $r, L_1(T)$ as in Theorem 2. The interaction between these invariants and the degree is governed by the Artin-Tate formula, which was used already to deduce Theorem 2 (by comparing (4.1) over \mathbb{F}_q and \mathbb{F}_{q^2}); here is the full statement, as in [7, Proposition 6].

THEOREM 8 (Artin-Tate formula). *If X satisfies the Tate conjecture,[†] then*

$$L_1(1) = |\Delta| \# \text{Br}(X), \quad (4.1)$$

where Δ denotes the discriminant of the Néron-Severi lattice and $\text{Br}(X)$ denotes the Brauer group; the latter is finite and its order is a perfect square. Also, the rank of the Néron-Severi lattice equals r ; in particular, if $r = 1$ then X admits a unique polarization, and its degree is equal to Δ .

We can now justify the choice of the conditions in Computation 6: besides smooth quartics, some other sources of degree-4 K3 surfaces include desingularization of singular quartics with only isolated rational singularities, which all have $r > 1$; and double covers of quadrics branched along $(4, 4)$ curves, which all have either $r > 1$, $L_{\text{alg}}(T) \neq 1 + T$, or $\Delta = -4$.

We now justify our previous assertion that Theorem 8 makes it difficult to produce enough examples to establish that Theorem 5 always holds for $q = 2$ with $n = 1$. Among the possibilities for $L_{\text{trc}}(T)$ allowed by Computation 3(c), there exist cases where

$$\deg L_{\text{trc}}(T) = 20, \quad L_{\text{trc}}(1) \in \{307, 367, 463\}, \quad L_{\text{trc}}(-1) = 3.$$

By Theorem 2, these can occur only with $L_{\text{alg}}(T) = 1 + T$. Since 307, 367, and 463 are prime, these would have to occur for K3 surfaces of degrees 2×307 , 2×367 , and 2×463 , respectively; however, the moduli spaces of polarized K3 surfaces of these degrees are of general type [9, Theorem 1], so constructing explicit points on them may be difficult. A more promising approach would be to make explicit the constructions used in [24], which involve lifting to characteristic 0, as in the proof of Honda's theorem.

[†]This is known to hold except in certain cases in characteristic 2; see [17]. The case of characteristic 2 is apparently resolved by a very recent preprint [15]. In any case, the formulation in [7, Proposition 6] is made carefully so as not to require the Tate conjecture.

Finally, we discuss the prospects for making similar calculations for $q > 2$. On the side of enumerating candidate zeta functions, there seems to be a bit of room to enlarge q ; for example, the following computation for $q = 3$ took 2.5 days (wall time, parallelized as in §2.3).

COMPUTATION 9. *The following sets are computed.*

- (a) *The set of polynomials $L(T)$ satisfying the conditions of Theorem 1 for $q = 3$; it contains 75,936,610 elements, representing 6,867,811 distinct values of $L_{\text{trc}}(T)$.*
- (b) *The set of polynomials in (a) consistent with Theorem 2; it contains 52,980,075 elements.*
- (c) *The set of polynomials in (b) consistent with the inequalities $\#X(\mathbb{F}_q) \geq 0$ and $\#X(\mathbb{F}_{q^{mn}}) \geq \#X(\mathbb{F}_{q^n})$ (it suffices to impose the second condition for $(mn, n) = (2, 1)$); it contains 49,645,728 elements.*

On the side of computing zeta functions of quartics, we must emphasize that the (optimized) brute-force approach we used to compute the zeta functions of all quartic surface over \mathbb{F}_q by point counting is not feasible for $q > 2$, particularly in the exceptional cases where one must go beyond $\mathbb{F}_{q^{11}}$ to resolve the sign ambiguity. For $q = 3$ there are already more than 2000 times as many PGL_4 -orbits of quartics to consider, and the time to compute $\#X(\mathbb{F}_{q^{11}})$ will be larger by a factor of at least 100 (more than 2000 for $\mathbb{F}_{q^{19}}$). One should instead look to methods based on p -adic cohomology, as in [1]; these have recently been made practical[†] for K3 surfaces [5].

Acknowledgements. This work was carried out at ICERM during the fall 2015 semester program “Computational aspects of the Langlands program”. We thank Edgar Costa, David Harvey, Brendan Hassett, Christian Liedtke, Sebastian Pancratz, Matthias Schütt, Lenny Taelman, and Yuri Tschinkel for helpful discussions.

References

1. T.G. ABBOTT, K.S. KEDLAYA, AND D. ROE, ‘Bounding Picard numbers of surfaces using p -adic cohomology’, in *Arithmetic, Geometry and Coding Theory (AGCT 2005)*, Séminaires et Congrès 21, Soc. Math. France, 2009, 125–159.
2. S. BASU, R. POLLACK, AND M.-F. ROY, *Algorithms in Real Algebraic Geometry*, second edition, Algorithms and Computation in Mathematics, 10, Springer-Verlag, Berlin, 2006.
3. E.R. BERLEKAMP, ‘Factoring polynomials over finite fields’, *Bell System Tech. J.* **46** (1967), 1853–1859.
4. H. COHEN, *A Course in Computational Algebraic Number Theory*, third printing, Springer-Verlag, Berlin, 1996.
5. E. COSTA AND Y. TSCHINKEL, ‘Variation of Néron-Severi ranks of reductions of K3 surfaces’, *Experimental Math.* **23** (2014), 475–481.
6. CYTHON – C-extensions for Python, <http://cython.org>.
7. A.-S. ELSSENHANS AND J. JAHNEL, ‘On Weil polynomials of K3 surfaces’, in *ANTS-XI: Algorithmic Number Theory Symposium*, Lecture Notes in Comp. Sci. 6197, Springer-Verlag, Berlin, 2010.
8. W. HART, F. JOHANSSON, AND S. PANCRAZ, FLINT: Fast Library for Number Theory, version 2.4.5, 2015, <http://flintlib.org>.
9. V.A. GRITSENKO, K. HULEK, AND G.K. SANKARAN, ‘The Kodaira dimension of the moduli of K3 surfaces’, *Invent. Math.* **169** (2007), 519–567.
10. S. HALOUI, ‘The characteristic polynomials of abelian varieties of dimension 3 over finite fields’, *J. Number Theory* **130** (2010), 2745–2752.
11. K.S. KEDLAYA, ‘Search techniques for root-unitary polynomials’, in *Computational Arithmetic Geometry*, Contemporary Math. 463, Amer. Math. Soc., 2008, 71–82. Associated code available at <http://kskedlaya.org/papers/>.
12. T. HONDA, ‘Isogeny classes of abelian varieties over finite fields’, *J. Math. Soc. Japan* **20** (1968), 83–95.
13. D. HUYBRECHTS, *Lectures on K3 Surfaces*, Cambridge Univ. Press, 2015.
14. M. LIEBLICH AND M. OLSSON, ‘Fourier-Mukai partners of K3-surfaces in positive characteristic’, *Ann. Sci. Éc. Norm. Supér.* **48** (2015), 1001–1033.

[†]The implementation used in [5] does not allow $q = 2$, so we cannot use it to check Computation 4.

15. W. KIM AND K. MADAPUSI PERA, ‘2-adic integral canonical models and the Tate conjecture in characteristic 2’, arXiv:1512.02540v1 (2015).
16. C. LIEDTKE AND Y. MATSUMOTO, ‘Good reduction of K3 surfaces’, arXiv:1411.4797v2 (2015).
17. K. MADAPUSI PERA, ‘The Tate conjecture for K3 surfaces in odd characteristic’, *Invent. Math.* **201** (2015), 625–668.
18. G. MALLE, ‘The totally real primitive number fields of discriminant at most 10^9 ’, Algorithmic number theory (ANTS-VII), Lecture Notes in Comput. Sci., 4076, Springer, Berlin, 2006, 114–123.
19. MANYPOINTS – table of curves with many points, <http://manypoints.org>.
20. J. MILNE AND W. WATERHOUSE, ‘Abelian varieties over finite fields’, in *1969 Number Theory Institute*, Proc. Symp. Pure Math. 20, Amer. Math. Soc., Providence, 1971, 53–64.
21. B. OSSERMAN, ‘The Weil conjectures’, in *The Princeton Companion to Mathematics*, Princeton Univ. Press, 2008, 729–732.
22. THE PARI GROUP, PARI/GP version 2.7.0, 2014, <http://pari.math.u-bordeaux.fr/>.
23. SAGE version 6.7, <http://www.sagemath.org/>.
24. L. TAEELMAN, ‘K3 surfaces over finite fields with given L-function’, arXiv:1507.08547v2 (2015); to appear in *Algebra and Number Theory*.
25. J. TATE, ‘Endomorphisms of abelian varieties over finite fields’, *Invent. Math.* **2** (1966), 134–144.
26. J. VOIGHT, ‘Enumeration of totally real number fields of bounded root discriminant’, Algorithmic Number Theory (ANTS-VIII), Lecture Notes in Comput. Sci. 5011, 268–281.
27. V. ZINOVIEV, ‘On the solutions of equations of degree’, Research Report RR-2829, INRIA, 1996, available at <http://hal.inria.fr/inria-00073862>.

Kiran S. Kedlaya
 Department of Mathematics
 University of California, San Diego
 9500 Gilman Drive #0112
 La Jolla, CA 92093
 United States of America
 kedlaya@ucsd.edu

Andrew V. Sutherland
 Department of Mathematics
 Massachusetts Institute of Technology
 77 Massachusetts Avenue
 Cambridge, MA 02139
 United States of America
 drew@math.mit.edu